Three Dimensional VLSI-Scale Interconnects

Dennis W. Prather

University of Delaware
Department of Electrical and Computer Engineering
Newark, DE 19716
email: dprather@ee.udel.edu

Abstract. As processor speeds rapidly approach the Giga-Hertz regime, the disparity between process time and memory access time plays an increasing role in the overall limitation of processor performance. In addition, limitations in interconnect density and bandwidth serve to exacerbate current bottlenecks, particularly as computer architectures continue to reduce in size. To address these issues, we propose a 3D architecture based on through-wafer vertical optical interconnects. To facilitate integration into the current manufacturing infrastructure, our system is monolithically fabricated in the Silicon substrate and preserves scale of integration by using meso-scopic diffractive optical elements (DOEs) for beam routing and fan-out. We believe that this architecture can alleviate the disparity between processor speeds and memory access times while increasing interconnect density by at least an order of magnitude. We are currently working to demonstrate a prototype system that consists of vertical cavity surface emitting lasers (VCSELs), diffractive optical elements, photodetectors, and processor-in-memory (PIM) units integrated on a single silicon substrate. To this end, we are currently refining our fabrication and design methods for the realization of meso-scopic DOEs and their integration with active devices. In this paper, we present our progress to date and demonstrate vertical data transmission using DOEs and discuss the application for our architecture, which is a multi-PIM (MPM) system.

Introduction

As modern day technologies continue to develop an increasing number of applications are resorting to computational based simulations as a tool for research and development. However, as simulation tools strive to incorporate more realistic properties their computational requirements quickly increase and in many cases surpass that which is currently available. As a result, a seemingly perpetual demand to process more information in shorter time frames has resulted. Moreover, while current computer architectures are steadily improving they are not keeping pace with the requirements of more sophisticated applications and in fact for some applications they are falling behind. To this end, new paradigm computer architectures need to be developed.

The current paradigm for addressing this shortcoming is to simply incorporate smaller devices into larger die. However, while this does enable the design and realization of more sophisticated circuits it also exacerbates an already serious problem, namely the interconnection and packaging of the devices and components within the system. For example, according to the National Technology Roadmap for Semiconductors, processors based on 1μm fabrication have a ratio of transistor -to-interconnect delay of 10:1 (assuming a 1mm long interconnect), whereas that for the same processor based on 0.1μm fabrication is 1:100. This represents a shift in emphasis of more than three orders of magnitude. As a result alternative interconnect and packaging technologies need to be developed. Therefore, in this paper we report on our work in addressing these technological barriers by designing an embedded processor-in-memory (PIM) architecture realized using an optically interconnected three-dimensional (3D) package.

While conventional 3D packaging increases circuit density, decreases interconnect delay, and reduces critical interconnect path lengths, their full potential has yet to be realized. This is due mainly to the associated capacitive and inductive loading affects of vertical vias, which reduce bandwidth and allow for only a 1-to-1 interconnect. To overcome these limitations we propose an alternate approach that is based on recent advances in micro-optical technology.

Our approach uses vertical cavity surface emitting lasers (VCSELs) that are flip-chip bonded onto CMOS drivers. The VCSELs have a 1.3μm wavelength which is transparent to the silicon wafer. The VCSELs are oriented such that the output beam is directed vertically through the silicon wafer. However, before the beam enters the wafer it is incident on a VLSI-scale diffractive optical element (DOE) that not only focuses the beam to a subsequent wafer, but also performs a 1-to-*N* fanout (*N* can range from 1 to 50 depending on the area used for the DOE). This allows for nearly real time data routing and distribution, which is essential to overcome conventional computational bottlenecks. However, before presenting further details in our approach we first motivate our PIM-based architecture.

PIM Motivation

A current trend in computer system design is to develop architectures based on the integration of a large number of smaller and more-simple processing cores that work together in unison. The idea here is that such processors can be integrated directly into random access memory (RAM) to simplify the memory hierarchy, i.e., level-1 and level-2 CACHE, and thereby streamline processor to memory communication. Such systems have been named Intelligent RAM (IRAM), Flexible RAM (Flex-RAM) and PIM, as we refer to it.

Currently, several high profile research initiatives (sponsored by federal agencies, e.g., the HTMT-PIM project [1,2], the DIVA project [3], the FlexRAM project [4] are investigating many of the architectural and system design issues related with the

implementation of PIM-based systems. In fact, IBM recently announced the introduction of the Blue-Gene project [5], which anticipates an industry investment on the order of a $100M dollars to produce a petaFLOPS scale machine based on thousands of PIM components. Therefore, even though PIM-based architectures are not currently being used in commercial machines, they promise to overcome the limitations of conventional computer architectures.

However, in general the amount of memory and the processing capability of individual PIMs is limited, therefore the construction of PIM-based high performance systems will require the integration of upwards of tens of thousands of PIMs. Thus the integration of multiple PIMs into a single package will be absolutely essential to reduce latencies, increase communication bandwidth between PIMs, reduce power consumption, and reduce the integration cost of the entire system.

Therefore the problem addressed by this research proposal is the implementation of a Multiple PIM Module (MPM) to harness the processing capability and the memory storage capability of multiple PIMs into a single computational module. A MPM can be used as the building block to implement mobile computers as proposed by the MIT RAW project. It can be used as the basic building block for computer systems specialized in data intensive computation, as proposed in the DIVA project. And it can be a building block for the DPIM region of a large scale, high performance computer such as the one proposed in the HTMT project.

Some of the open research problems in the implementation an MPM and in its use in a system architecture are: (1) How the multiple PIMs, that form the MPM, communicate and synchronize with each other. (2) Is it possible to design and implement a fast and versatile interconnection between the multiple PIMs in the MPM. (3) How MPMs can be programmed and how the interconnection can be adapted for new communication pathways. And (4) how does the runtime system control MPMs to ensure the communications/synchronizations are performed in the most efficient way according to the needs of the application program.

To address these issues, we are developing a technology based on the interconnection of multiple PIMs within a single MPM via an array of vertical cavity surface emitting laser (VCSEL) and SiGe detector arrays that are vertically interconnected through the silicon wafer using a DOE. This technology allows for fast, abundant, and distributed interconnections amongst the PIMs in a given module. Also, because this approach allows for data distribution at the 2-5GHz rate it reduces the latency in communication between PIMs to unprecedented levels and because optical beams can essentially pass right through each other without exchanging information it all but eliminates the place and route problem. Also, each interconnect link in our design would consume approximately 50mW of power, which when applied to a full 16×16 interconnection would consume on the order of 10 Watts of power. This is nearly an order of magnitude less than current architectures that are limited to only 4×4 interconnections.

To realize this architecture three critical technologies must be used: long wavelength VCSELs (1.3µm), high speed (2-5 Gbits) CMOS drivers for the VCSELs, and VLSI-scale DOEs. To this end we have been working with Gore Photonics for the 1.3µm VCSELs and developing our own high-speed CMOS drivers, VLSI-scale DOEs, and system integration techniques at UD. Thus, in the remainder of the paper we report on our progress in this effort. We begin by motivating the 3D architecture and the describe the component optical technologies needed to realize it.

Optoelectronic Technologies

Whereas the use of optical interconnects in long haul and local area networks has proven extremely successful, its use on the VLSI-scale has been limited. This is due in large part to the continual increase in speed and performance of conventional electronic devices. However, the issues associated with next generation PIM architectures cannot be adequately addressed with speed alone. Instead, such systems will require not only the ability to share or distribute information among PIM modules (signal fan-out) but also a significant increase in interconnect density. While the issues of increased bandwidth, interconnect densities, and signal fan-out are individually compelling reasons for considering optical interconnects, when combined together they become persuasive. For example, one possible electronic solution to increasing interconnect density is to use flip-chip, or bump, bonds, which can require approximately 20µm^2 of chip area while offering only a 1-to-1 interconnect. In comparison, we have designed VLSI-scale diffractive optical elements (DOEs) that within the same area provide a 1-to-16 interconnect. Currently we have experimentally demonstrated a 1-to-4 and are in the process of fabricating the 1-to-16. For this reason we propose the use of an optoelectronic 3D architecture that uses monolithically integrated VLSI-scale DOEs for application to PIM architecture, as shown in Fig. 1.

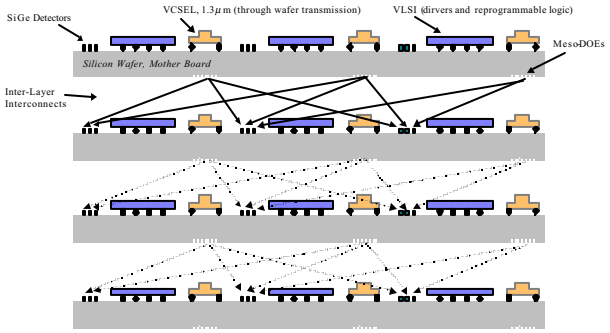

Fig. 1. Monolithic interconnect architecutre that uses 3D diffractive optical interconnects on the VLSI-scale for through wafer fan-out interconnects for data or clock distrbution. Various modules of this architecture can also be stacked together to realize more complex systems.

While the notion of 3D architectures is appealing, due to the efficient use of power and increased processing and interconnect densities, few of the systems

proposed in the literature have received wide spread use. Reasons for this depend on the technology being used. For instance, all electronic architectures suffer from either reduced communication bandwidth, due to routing the inter-layer interconnects through the periphery of the 3D stack, or reduced interconnection density, due to the inability to distribute data between layers using 1-to-1 bump bonds. Along the same lines, optical architectures suffer from input/output coupling efficiencies, for wave guide based approaches, interconnect density and distribution, for 1-to-1 emitter-receiver-based approaches, and scale of integration for bulk optical systems. Thus, we believe that in order for an optical interconnect system to be viable it *must* satisfy the following conditions: (1) It must have a scale of integration comparable to VLSI, to preserve scales of integration. (2) The optical system must be monolithic in the Silicon substrate, in order to alleviate alignment issues and improve system reliability. And (3) the fabrication methods and materials used must be compatible with the current manufacturing infrastructure, in order to reduce cost of implementation. In the design of our architecture we will are strictly adhering to these conditions.

Our approach is based on our recent progress in the development of suitable design tools, which enable the design of VLSI-scale DOEs for monolithic integration with active devices. As a result, we have been able to significantly increase the interconnection density as compared to all electronic vertical interconnections as illustrated in Fig. 2, which illustrates a DOE that occupies $10\mu m^2$ and provides a 1-to-4 fan-out. If this DOE is tiled over a $20\mu m^2$ area, equivalent to that of a bump bond, it would provide a 1-to-16 fan-out in comparison to a 1-to-1, which represents more than an order of magnitude increase in interconnect density. In addition to increasing density this approach significantly simplifies the place and route problem because optical beams do not exchange information and can therefore accommodate overlap in the routing process.

In order to realize optical interconnections within a Silicon wafer and on a scale comparable with VLSI circuits, one must be able to heterogeneously integrate active and passive optical devices together on a scale comparable to microelectronic devices. This must also be done in such a way that the ability to control and redirect light in a general fashion is preserved, e.g., off-axis focusing, mode shaping, and beam fan-out. Whereas active optical devices, such as emitters, detectors, and modulators are readily designed and fabricated with dimensions on the micron scale, until recently passive optical elements capable of such general behavior were not. However, recent advances in both the design and fabrication of diffractive structures [6] now enables the integration of active and passive optical devices on the VLSI-scale and the ability to efficiently control and redirect light in a general fashion, see Fig. 3. Thus, the integration of VCSELs with wavelength scale fan-out DOEs on the VLSI-scale offer not only an order of magnitude improvement (in terms of density, bandwidth, and power consumption) but also the ability to design architectures that heretofore have not been possible. As a result new optical interconnect architectures can now be developed.

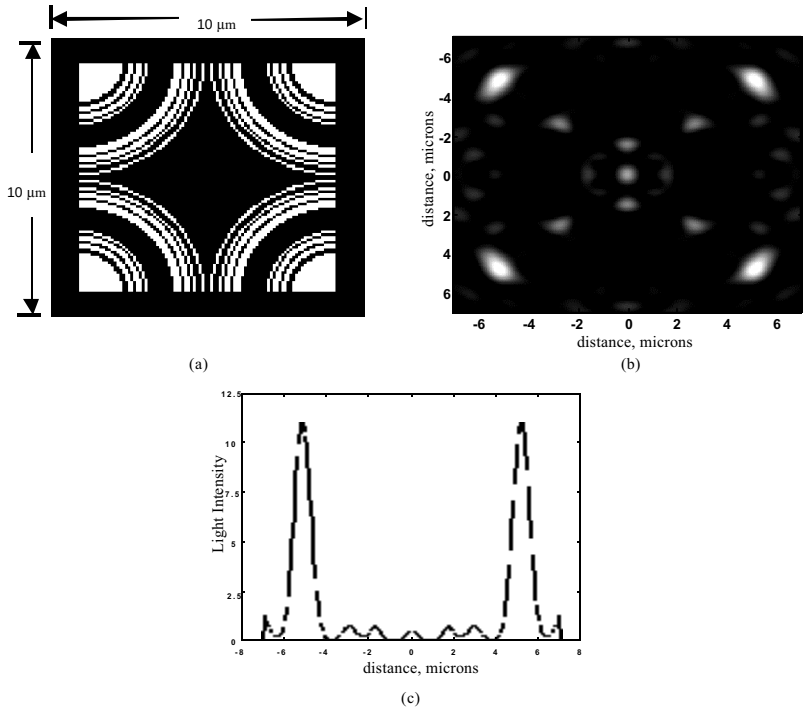

Fig. 2. Illustration of a three dimensional subwavelength off-axis lenslet array used for 1-to-4 fanout on the VLSI-scale (a) DOE, (b) intensity image in the focal plane, and (c) line scan thorough the focal plane. Results were generated using a 3D FDTD diffraction model.

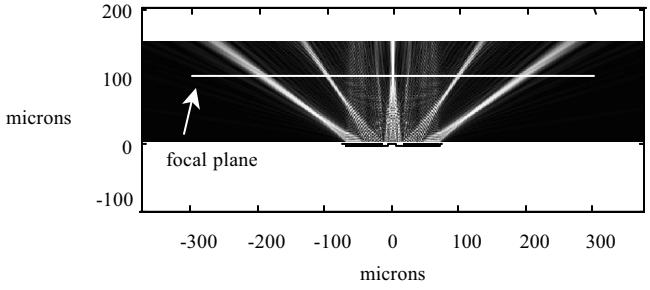

Fig. 3. Illustration of a VLSI-scale 1-to-5 fan-out DOE, computed using the boundary element method. The width of the DOE is 120 microns and the focal length is 100 microns.

Recently we have fabricated and experimentally validated these elements and are currently preparing them for system level integration [7]. However, critical to the successful completion of this effort is the ability to fabricate DOEs that have features sizes on the nanometer scale. Although many fabrication techniques for DOEs exist, by far the most general and widely used is that of the microelectronics photolithographic process. In this technique the profile of a DOE is realized by etching micro-relief patterns into the surface of either conducting or dielectric substrates. A curved surface profile is realized by using a multi-step process which produces a stair-step approximation. Using this fabrication process DOEs that have diffraction efficiencies on the order of 95% have been fabricated. Unfortunately, as the scale of a DOE is reduced the alignment process, needed for multi-step profiles, becomes exceedingly difficult. As a result, alternate fabrication methods based on single step gray-scale lithography and direct electron beam (e-beam) exposure have been developed.

In the gray-scale process one wishes to realize continuous profiles, or structures. However, for devices on the VLSI-scale current fabrication technology limits us to a discrete number of levels, typically 4-8 levels. Thus, we can currently fabricate our DOEs using a gray-scale technique which results in multilevel structures from a single processing step, as shown in Fig. 4. To this end, we designed our multi-level masks in the lab and used an outside vendor [8] to provide the gray-scale mask. Once we have the mask we deposit an initial height of the photoresist on the substrate, i.e., silicon wafer, which can be precisely controlled by adjusting the spin rate at the time of deposition. Through experimentation, we have characterized the response of photoresist to various degrees of UV exposure. This allows us to precisely designate the correct transmission levels of the mask to create our multi-level DOE profiles in the photoresist.

After the grayscale photolithography, the pattern is transferred into the surface of the Silicon substrate using a Plasmatherm 790 series reactive ion etching (RIE) system. Careful calibration of the RIE process is required to achieve structures with smooth surfaces and submicron feature resolution while preserving the height of the initial profile.

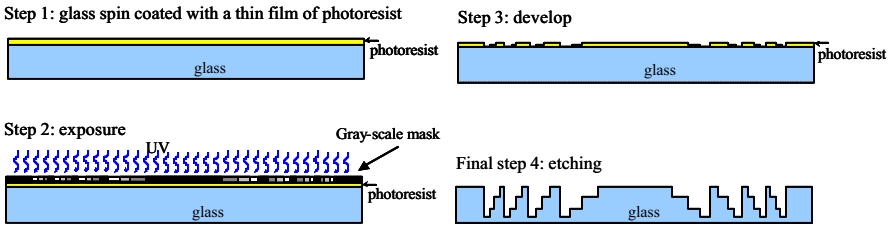

Fig. 4. Graphical illustration of the gray-scale photolithographic fabrication process.

An alternative fabrication method based on direcet e-beam write can also be used to fabricate VLSI-scale DOEs. In this appraoch a high energy electron beam is

used to expose a photoresist coated substrate. As the substrate is exposed the energy level of the e-beam is varied in accordance with the desired DOE profile. Once developed the substrate is etched, using techniques such as reactive ion etching, to transfer the continuous photoresist profile into the substrate, see Fig. 5. This process is capable of fabricating binary DOE profiles that have feature sizes on the order of 60nm, which is several times smaller than the wavelength of illumination. As a result efficiencies exceeding those predicted by scalar diffraction theory can be achieved [9]. Through collaboration with Axel Scherer of CalTech we have recently had several DOEs fabricated, as shown in Fig. 6.

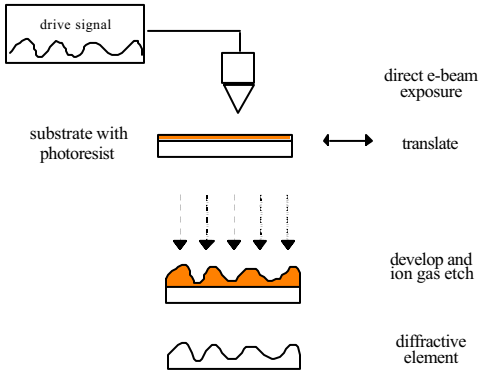

Fig. 5. Fabrication process for continuous profile DOEs based on direct electron beam write.

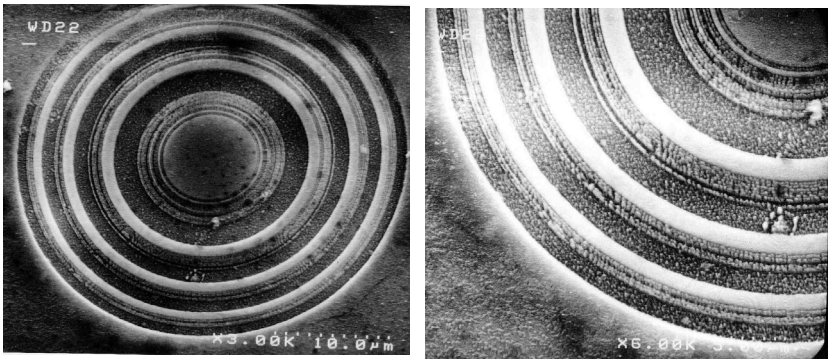

Fig. 6. Illustration of a mesoscopic diffractive lens having a diameter of 36µm, a focal length of 65µm and a minimum feature size of 60nm. The element was fabricated by Dr. Axel Scherer, of the California Institute of Technology.

In addition to developing the theoretical and experimental framework necessary to design and realize DOEs we have developed a novel system for characterizing their performance.

Our system consists of a microscope objective (20X) and a 1inch diameter lens. The system has an overall magnification of 4.2 (based on the ratio of the two focal lengths, f_2 / f_1), and is able to resolve 1 micron minimum features. The entire imaging system is mounted on an *x,z* translation stage, as shown in Fig. 7. Because the object and image planes, in this system, are fixed and well defined they can be used to determine the axial location relative to the DOE, i.e., the reference plane for z=0. This is achieved by translating the imaging system toward the DOE until the surface is imaged on to the CCD. Subsequently, the translation stage, with the entire imaging system on it, is translated back to the plane of interest, i.e., $z=z_o$. Because the microscope objectives have large numerical apertures the performance of the imaging system, i.e., its modulation transfer function (MTF), reproduces the intensity profile in the object plane, i.e, the observation plane, with excellent fidelity.

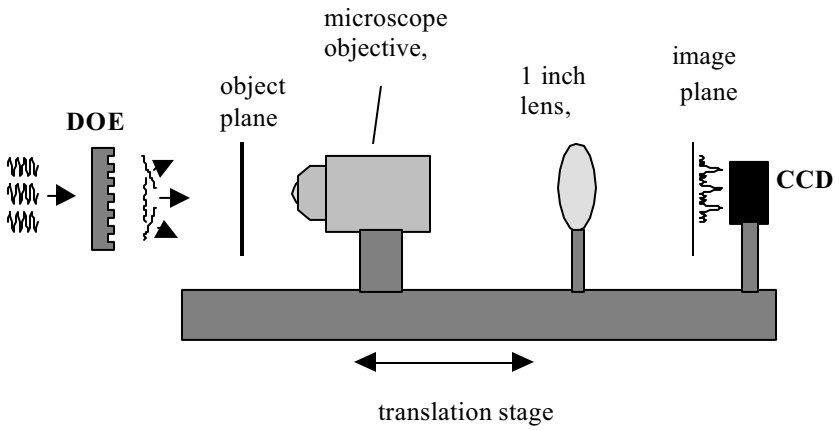

Fig. 7. Micro 4f imaging system for characterizing mesoscopic diffractive optical elements.

To validate our electromagnetic design models we used the system to measure the diffracted light from a precision pin-hole of 71μm in diameter, from a collimated incident wave of 0.633nm. We then calculated the diffracted light using both scalar diffraction theory and using our electromagnetic model, results for z=350μm are shown in Fig. 8. Additional measurements were made along the *z*-axis and showed the same level of agreement. To illustrate the utility of this system we used it to characterize the diffractive lens shown in Fig. 6, the results are shown in Fig. 9. Once confident that our design and fabrication methods were working we then applied them to the realization of through silicon wafer DOEs [7].

Integration

In order to achieve optical interconnects on a single Silicon die, we must be able to integrate emitters, detectors, drivers, and DOEs on the VLSI scale. Our

approach toward integration will be to construct a hybrid system using flip-chip bonding. For this part of the project we will use a SEC Omnibonder 860 flip-chip bonding machine to construct a multichip module for the integration of the active and passive optical devices with their electronic counterparts. Figure illustrates the integration of an 8×8 CMOS driver array with an 8×8 980nm VCSEL array.

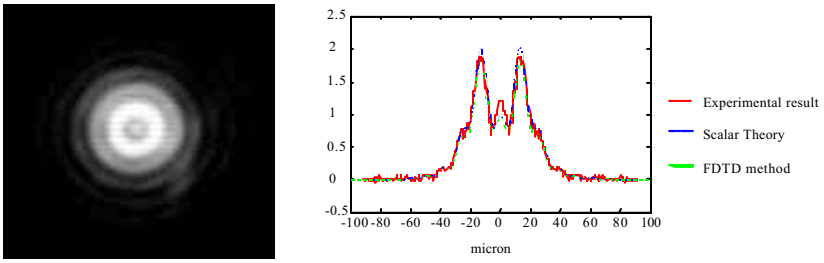

Fig. 8. Comparison between experimental results and theoretical predictions for the diffraction from a precision pin-hole that had a diameter of 71 microns at an axial location of 350 microns.

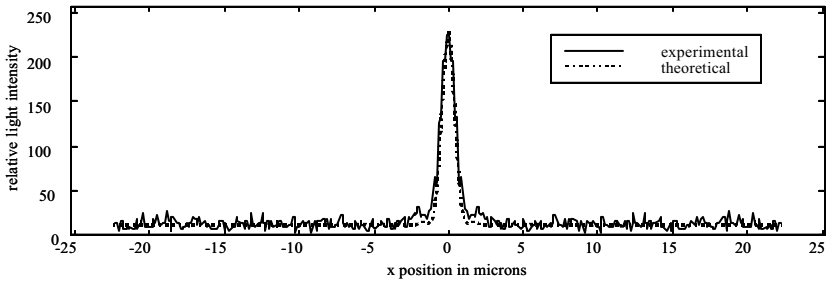

Fig. 9. Overlay of the experimental characterization of a mesoscopic diffractive lens and the results predicted from our electromagnetic models. Data was taken from our system using a 40X magnification objective at the loacation of z=65µm, the design focal length.

Ultimately, we plan to use 1.3 micron VCSELs as emitters and a Silicon substrate as the medium of propagation. However, such long wavelength VCSELs are not currently available in die form at present, so we have begun the construction of a pre-prototype system using an 850nm and 980nm VCSELs on a glass substrate. In this preliminary system, the VCSEL is bonded to a CMOS driver circuit and directed through the DOE as shown in Fig. 10.

Our main concern associated with bonding the VCSEL over a DOE, is the air gap spacing between the VCSEL and the backside of the glass substrate. Since the VCSEL will be flip-chip bonded to the glass surface, the solder bump size, bond pressure and bond temperature profile will affect the resultant air gap. Additionally, the proximity of the CMOS driver and the VCSEL will be a guiding parameter of the

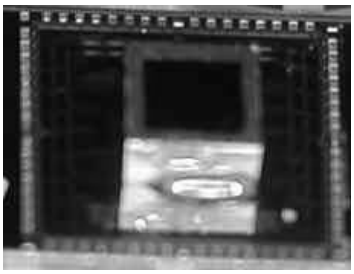

Fig. 10. Illustration of a VCSEL flipchip bonded to a CMOS driver circuit. The VCSELs and CMOS drivers were supplied by the U.S. Army Research Laboratory.

bonding temperature profile, since we do not want the first device bonded to detach during the second bond. Most likely, we will choose to bond the CMOS driver first in order to maximize control over the air gap spacing. That way heating during the bonding of the driver will not affect the final VCSEL position.

Summary

We have discussed the motivation for chip-level optical interconnects, and proposed a 3D architecture that offers higher bandwidth interconnect density in comparison to conventional architectures. Also, we have discussed a potential applications for our architecture based on a multi-processor-in-memory system. To this end, we demonstrated through-wafer optical fan-out using VLSI-scale DOEs and long wavelength VCSELs (courtesy of Gore Photonics). Flip-chip bonding gives us the ability to integrate active and passive devices on a single die, and we are currently building a prototype system to demonstrate this integration. The significance of our approach lies in the ability to design optical elements that efficiently control, or redirect, light on a VLSI-scale and can be directly integrated into the current VLSI-based manufacturing infrastructure. As such this technology lends itself nicely to 3D interconnect schemes and facilitates the trend toward higher levels of parallelism in computer architectures.

References

[1] T. Sterling, “Achieving petaflops-scale performance through a synthesis of advanced device technologies and adaptive latency tolerant architectures,” in *Supercomputing 99*, (Portland, OR), Novermber 1999.

[2] P.M. Kogge, J.B. Brockman, T. Sterling, and G. Gao, "Processing-in-memory: chips to petaflops," in *International Symposium on Computer Architecture*, (Denver, CO), June 1997.

[3] M. Hall, P. Kogge, J. Koller, P. Diniz, J. Chame, J. Draper, J. LaCoss, J. Granacki, A. Srivastava, W. Athas, J. Brockman, V. Freeh, J. Park, and J. Shin, "Mapping irregular applications to DIVA, a PIM-based data-intensive architecture," in *Supercomputing 99*, Portland OR, Novermber 1999.

[4] Y. Kang, M. Huang, S.M. Yoo, Z.Ge, D. Keen, V. Lam, P. Pattnaik, and J. Torrellas, "Flexram: toward an advanced intelligent memory system," in International Conference on Computer Design, October 1999.

[5] IBM, "IBM unveils $100 million research initiative to build world's fastest Semiseek, December 1999.

[6] D.W. Prather, M.S. Mirotznik, and S. Shi, *Mathematical Modeling in Optical Science*, Ch. Electromagnetic models for finite aperiodic diffractive optical elements, *in print*, SIAM Frontier Book Series, Society for Industrial and Applied Mathematics, 2000.

[7] M. LeCompte, X. Gao, H. Bates, J. Meckle, S. Shi, and D.W. Prather, Three-dimensional through-wafer fan-out interconnects," in Optoelectronics Interconnects VII, SPIE 3952, The International Society Optical Engineering, Bellingham WA, January 2000.

[8] Canyon Materials, Inc., San Diego, CA.

[9] J.N.Mait, D.W. Prather, and M.S. Mirotznik, "Binary subwavelength diffractive-lens design," *Opt. Lett*., **23,** pp. 1343-1345, September 1998.